

Frequently asked questions

Where is SQL Workload Analysis installed? Should I install anything on the database machine?

SQL Workload Analysis monitors SQL Server instances remotely (agent-less), from a dedicated framework machine. SQL Workload Analysis has a web-based interface that you can use, via a browser, from any machine that can access the framework machine.

How is SQL Workload Analysis different from the Precise for SQL Server Enterprise edition?

It is based on the same collection technology, but there's a completely new UI, and it is designed from the bottoms up to be extremely easy to install and to use. SQL Workload Analysis is focused on SQL Server performance tuning, stemming from 20 years' R&D and DB tuning experience - all packaged in a user-friendly and effortlessly-deployed application. Precise for SQL Server Enterprise Edition complements SQL tuning with full-blown enterprise functionality including alerts, reports, up-stream application context, and network monitoring.

How can I pause monitoring an instance?

When adding monitoring for an instance, a dedicated Windows Service is created on the framework machine. Stopping/starting this Service will stop/start the monitoring for that instance.

I see a performance peak in one of the graphs. How do I zoom in to investigate it further?

SQL Workload Analysis has an easy zoom functionality. Simply drag your mouse with the left button click over the graph, drawing a rectangle to define the area you want to zoom in on. SQL Workload Analysis will zoom in quickly on this area. Note: zoom is enabled in all graphs besides in "Current" pane. Graphs displaying current information always show last 30 minutes of data in maximal granularity.

Data has stopped being collected. How do I troubleshoot this?

Troubleshoot data collection pauses with the following: Check whether the Windows Service of the collector is up and running. If it is stopped, restart it. Check whether the monitored instance's credentials were changed. If this is the case, SQL Workload Analysis will not be able to connect to the monitored instance until the new credentials are provided. Rerun Setup.exe from the installation root folder to update credentials.

How to monitor in failover cluster configuration?

Install the SQL Workload Analysis framework on a dedicated server outside of the cluster. During setup, specify the logical (cluster) instance name as the SQL Server instance to be monitored. In case of failover, the collector will identify the failover in a second, and within moments will resume monitoring automatically without any manual intervention.

How do I upgrade from trial to standard version to continue monitoring beyond the trial's period?

Visit the [IDERA Online Store](#) to purchase a license.

How can I control the time frame of the displayed data?

SQL Workload Analysis provides simple shortcuts for setting desired time frame (e.g. last 8 hours) – these shortcuts are displayed on the top right part of the screen. Besides the shortcuts you can always zoom in from any graph (except for "Current" pane that always shows last 30 minutes).

What does "Total in DB" measure?

"Total in DB" is the total amount of time spent by the application/s inside the database over the selected time period. This measurement allows you to quickly identify the activities with the most significant contribution to DB resources usage. This time is divided into productive "Using CPU" time and non-productive waits for resources. This breakdown is clearly displayed along with the absolute time measurements.

What is the refresh rate of the SQL Workload Analysis views?

The main dashboard portlets are refreshed every 15 minutes. The "Current" pane's portlets are refreshed every minute. The rest of the screens are not automatically refreshed since they are designed for investigating performance bottlenecks and should not be refreshed without user's intention.

What are the units I see next to the time displays (including in the Y axis of the graphs)?

The units are: μ s – microsecond, ms – millisecond, s – second, m – minute, h – hour, d – day, w – week.

For some SQL statements, why don't I see summary values regarding physical reads, logical reads, etc.?

These values are not accessible whenever a statement is not kept in the SQL Server's procedure cache.

What are the "Top Accessed Objects" I see for a SQL statement?

Accessed objects are the tables and the indexes the statement is utilizing during its execution. This display allows you to understand object bottlenecks even without analyzing the execution plan.

How should I work with the Locking Sessions (Blockers) portlet to analyze locking situations?

The sessions appearing on the top of the hierarchy are the blocker sessions. They are also colored in red. To release the lock you may need to consider killing these blocker sessions. The number that you can see after the program's name is the client process ID, which should help you in locating the blocker. The sessions appearing beneath the blocker sessions are the locked sessions. Each locked session can in its turn cause additional sessions to become locked – these locked sessions will appear further down in the hierarchy. When looking at the locked session this column presents the locked statement. When looking at the blocker session this column presents the last statement that was executed by the session. Many times it will represent the locking statement.

The main dashboard shows only top databases/statements. How can I see more in the list?

Drill down through the links provided inside the lists to get to the designated page showing much longer lists and the important details for each item.

What does the "Table Cardinality" column mean in the statement's execution plan's display?

This column shows the number of rows in the access table. It may appear when the Operator name is: Table scan, Index seek, or Clustered index seek. Table cardinality should be used in conjunction with the value of Returned row column. Returned rows shows how many rows were returned after applying all the filter predicates for the selected operation. It allows calculating filtering ratio and assessing the potential of using an index. Let's assume: Operation name is Table scan, Table cardinality is 1000000, and Returned rows is 50. It means that SQL Server sequentially scans a 1000000-row table while only 50 rows were qualified after applying relevant filter predicates. Normally an index may help when the filter ratio (Returned rows / Table cardinality) is less than 1%. You may view filter predicates applied by an Operation by hovering on the Operation and then clicking the More properties link in the presented tool tip.

Locks trending graph shows lock types including "MD" prefix – what is it?

"MD" stands for METADATA. It means that the locked object is not user data such as table row or index key. The locked object is some dictionary object that describes user objects. In most cases METADATA locks are caused by DDL (create, alter, drop) rather than DML (insert, update, delete) statements. "MD Partition Lock" means that the partition function of a partitioned table is locked. It may happen when adding / removing partitions. During that time, any application trying to access that partitioned table will be locked. "MD Statistics Lock" means that the lock is due to trying to gather statistics on user objects and update the statistics in system tables. Updating stats in system tables may lock the optimizer when trying to generate an execution plan for queries that access objects on which statistics were gathered. "MD Other Lock" stands for any other METADATA lock (not one of the two mentioned above). The full list of METADATA sub types reported by SQL Server is documented as part of column resource_subtype of table sys.dm_tran_locks.

How can I extend SQL Workload Analysis' visibility into EMC storage?

SQL Workload Analysis provides deep visibility past the database and into the EMC storage array, including the ability to detect contention on physical disks or virtual pools. In order to enable this extended visibility, some EMC software needs to be installed on the SQL Workload Analysis server as well as the database hosts. Please consult the following document for the full details of how to extend SQL Workload Analysis' visibility into EMC Storage: [SQL Workload Analysis – Installing and Verifying EMC Solution Enabler](#).

Why do certain statements show zero execution counts?

IDERA SQL Workload Analysis computes the top SQL statements in terms of time spent in the database from one source of such data in the system tables of SQL Server and correlates this with the execution count of that statement from another source (i.e another set of system tables in the SQL Server). Because of sampling rates, sampling errors, and the transient nature of the latter source, the correlation is sometimes not possible for certain statements with very few executions in any given time frame.